

Budget-Balanced and Nearly Efficient Randomized Mechanisms: Public Goods and Beyond*

Mingyu Guo¹, Victor Naroditskiy², Vincent Conitzer³, Amy Greenwald⁴, and Nicholas R. Jennings²

¹ Computer Science Department, University of Liverpool, UK
Mingyu.Guo@liverpool.ac.uk

² School of Electronics and Computer Science, University of Southampton, UK
{vn,nrj}@ecs.soton.ac.uk

³ Department of Computer Science, Duke University, USA
conitzer@cs.duke.edu

⁴ Department of Computer Science, Brown University, USA
amy@cs.brown.edu

Abstract. Many scenarios where participants hold private information require payments to encourage truthful revelation. Some of these scenarios have no natural *residual claimant* who would absorb the budget surplus or cover the deficit. Faltings [7] proposed the idea of excluding one agent uniformly at random and making him the residual claimant. Based on this idea, we propose two classes of public good mechanisms and derive optimal ones within each class: Faltings’ mechanism is optimal in one of the classes. We then move on to general mechanism design settings, where we prove guarantees on the social welfare achieved by Faltings’ mechanism. Finally, we analyze a modification of the mechanism where budget balance is achieved without designating any agent as the residual claimant.

1 Introduction

Many scenarios where participants hold private information require payments to encourage truthful revelation. Some of these scenarios have no natural *residual claimant* who would absorb the budget surplus or cover the deficit (e.g., a group of roommates deciding who gets to use the living room for a weekend party or a company distributing free football tickets among employees). Mechanisms with budget deficit are not very compelling as they require a subsidy. In more compelling surplus-generating (or, weakly budget-balanced) mechanisms, the surplus represents a loss in social welfare (i.e., the sum of the agents’ utilities), which can be viewed as the cost of truthfulness. A number of recent papers have investigated what the minimum budget surplus is that still supports truthful reporting and efficient outcomes [14, 11, 12, 4, 1, 2]. While weak

* Part of the work was performed while the first four authors were at Centrum Wiskunde & Informatica (CWI) in Amsterdam. Conitzer gratefully acknowledges NSF IIS-0812113, IIS-0953756, and CCF-1101659, as well as an Alfred P. Sloan fellowship, for support. Naroditskiy and Jennings gratefully acknowledge funding from the UK Research Council for project ‘Orchid’, grant EP/I011587/1. Greenwald gratefully acknowledges NSF Grant CCF-0905234.

budget balance is a necessary assumption,⁵ efficiency is not. In fact, sacrificing efficiency leads to a higher social welfare in certain cases (by having significantly lower net payments than efficient mechanisms) [7, 10, 5].

The mechanisms we propose here are budget-balanced (i.e., no loss of social welfare is due to the budget surplus) but not efficient. Our work starts with the idea behind Faltings' mechanism [7], which is that we exclude one agent uniformly at random and make him the residual claimant of the payments collected by an efficient mechanism (e.g., the VCG mechanism) in the market with only the remaining agents. Crucially, in order to maintain truthfulness, the outcome must be chosen without considering the private value of the excluded agent. Thus, the chosen outcome may not be the same as the efficient outcome when all agents' values are considered. This results in a social welfare below the value of the efficient outcome. Notice that the loss of social welfare is due only to the non-efficiency of the outcome as mechanisms with a residual claimant are budget-balanced. Since excluding one agent at random results in a randomized outcome function, we speak of expected social welfare. We say that a mechanism is r -competitive if its expected social welfare is at least r of the value of the efficient outcome for all types the agents may have: i.e., we are using a worst-case metric.

We apply the approach of excluding one agent to the public project scenario where a group of agents needs to decide whether or not to build, say, a bridge that comes at a publicly known cost. The public project scenario is fundamental to mechanism design: unlike allocation scenarios, no agent can be excluded from enjoying the benefits of the project if it is undertaken. While maximizing social welfare has been studied extensively in allocation scenarios (see e.g., [14, 11, 9, 10, 5]), public good scenarios received relatively less attention. [1] and [2] both studied the problem of designing welfare-maximizing public good mechanisms. [1] studied one dominance relationship between mechanisms, but did not propose any specific mechanisms. [2] studied sequential public good mechanisms with a different notion of truthfulness. Then, there are several general mechanisms that can be applied to public project [3, 4, 7]. The mechanism described in [3] is not budget-balanced when applied to public project, while the mechanism [4] has a zero competitive ratio. The paper by Faltings [7] is central to many of our results as we discuss next.

First, we derive a competitive ratio for the mechanism proposed by Faltings and prove its optimality within a class of mechanisms. Specifically, we define a class of mechanisms based on the fraction of the cost of the project passed on to the excluded agent. It turns out all mechanisms that assume the excluded agent would cover up to $\frac{1}{n-1}$ of the cost are $\frac{n-1}{n}$ -competitive (n is the number of agents). Faltings' mechanism corresponds to the excluded agent covering $\frac{1}{n}$ of the cost and is optimal within the class. Mechanisms that assume the excluded agent covers more than $\frac{1}{n-1}$ of the cost are not $\frac{n-1}{n}$ -competitive. A natural question is whether a better mechanism is possible. To this end, we consider a larger mechanism class by taking the mixtures over the above mechanisms. That is, we consider mechanisms that assume the excluded agent would cover a randomized proportion of the cost. We characterize one optimal mechanism within this larger class, which turns out to be $\frac{n}{n+1}$ -competitive.

⁵ An arbitrary social welfare can be achieved when unlimited subsidies are allowed.

The mechanisms above make the excluded agent the recipient of the VCG payments computed without him. The idea of computing VCG payments after excluding one agent has also been used in general quasi-linear domains to design redistribution of VCG payments computed with all agents present. Specifically, in a regular VCG mechanism, the rebate to agent i can be set to $\frac{1}{n}$ of the VCG payments collected in the market without him [3]. The resulting mechanism is efficient but not budget-balanced, and may run a deficit in the public good scenario.

We find that using this redistribution idea together with the inefficient allocation, made after excluding one agent, leads to a budget-balanced mechanism that *does not designate any agent as the residual claimant*. This results in a more fair treatment of all agents, and we call the mechanism *FaltingsFair*. In more detail, we set each agent's payment to be the expected VCG payment he would make after one of the other agents is excluded uniformly at random. This payment is reduced by the rebate described above. The sum of the rebates cancels out the sum of the payments, thus achieving budget balance. Interestingly, this mechanism was already proposed by Faltings in extended versions of his work [6, 8], though without the redistribution interpretation. Our analysis sheds new light on this mechanism establishing connections to a standard redistribution function and providing novel proofs.

The rest of this paper is structured as follows. A general model of mechanism design problems is stated in Section 2. Mechanisms with a residual claimant for the public good scenario are studied in Section 3. There we propose two classes of mechanisms and derive optimal ones within each class. In Section 4, we move on to general mechanism design settings. One of the optimal public good mechanisms we derive in Section 3 turns out to be a special case of Faltings' mechanism. We modify this mechanism to remove the residual claimant, which results in the budget-balanced *FaltingsFair* mechanism. Discussion of the results appears in Section 5.

2 Model

The set of agents is denoted by N ($|N| \geq 3$) and the private type of agent $i \in N$ is given by θ_i . The mechanism chooses an outcome $k(\theta')$ from the set of possible outcomes K , based on the profile of reported types θ' . The value of an agent for each outcome depends on his type $v_i(k(\theta'), \theta_i)$, and the utility is quasi-linear. Given an outcome $k \in K$ and a payment $t_i \in \mathbb{R}$, the utility is $u_i(k, t_i, \theta_i) = v_i(k, \theta_i) - t_i$. Let $k^*(\theta)$ denote the *efficient* outcome $k^*(\theta) \in \arg \max_{k' \in K} \sum_i v_i(k', \theta_i)$.

The VCG (also known as Clarke or pivotal) mechanism is defined by the efficient outcome and the following payments *from* the agents: $t_i^{\text{vcg}}(\theta) = \sum_{j \neq i} v_j(k^*(\theta_{-i}), \theta_j) - \sum_{j \neq i} v_j(k^*(\theta), \theta_j)$, where $k^*(\theta_{-i}) \in \arg \max_{k' \in K} \sum_{j \neq i} v_j(k', \theta_j)$.

3 Public Project

In a public project (equivalently, public good) problem, a group of agents needs to decide whether or not to undertake a project such as building a bridge. The two possible outcomes are: do not build the bridge and distribute C among the agents or build the bridge spending C on its construction. Each agent has a private value θ_i for having the

bridge built. We define the value of the efficient outcome as $\max(\theta_N, C)$: the sum of agents' values is $\theta_N = \sum_{i \in N} \theta_i$ when the bridge is built and C when it is not built. The valuation function of agent i consistent with this definition of social welfare is

$$v_i(k(\theta), \theta_i) = \begin{cases} \theta_i & \text{if } k(\theta) = 1 \\ \frac{C}{n} & \text{otherwise} \end{cases} \quad (1)$$

Faltings' mechanism [7] is defined as follows (we will call it *Faltings* from now on):

- We exclude one agent uniformly at random.
- The remaining agents use the VCG mechanism to come up with an optimal allocation for themselves.
- The excluded agent acts as the residual claimant. That is, the VCG payments are redistributed to the excluded agent, to achieve budget balance.

Faltings is known to be (dominant-strategy) incentive compatible and budget-balanced.⁶ *Faltings* can be generalized to the following class of mechanisms (also incentive compatible and budget-balanced):

- We pick one agent, denoted by a , uniformly at random, and we pretend agent a 's reported type is $C - x$ (ignoring what a actually reported).
- All agents, **including** a , participate in a VCG mechanism.
- a acts as the residual claimant. That is, everyone excluding a pays his VCG payment to a . (a does not have to make any payment. Note that incentive compatibility for a is guaranteed because a 's report is ignored altogether.)

Mechanisms inside the above class are characterized by the parameter x , where x represents how much the non-excluded agents need to value the project in order for it to be built. When there is no ambiguity, we will simply use mechanism x to refer to the mechanism inside the class that is characterized by x . *Faltings* corresponds to $x = \frac{n-1}{n}C$: For this value of x , the decision is to build if and only if the remaining agents' total valuation is at least $\frac{n-1}{n}C$, which is efficient for the remaining agents.

The parameter x could take any value in $(-\infty, \infty)$, but we only need to consider $x \in [0, C]$ (assuming non-negative types). We recall that x represents how much the non-excluded agents need to value the project in order for it to be built. When $x < 0$, mechanism x is equivalent to mechanism $x = 0$ in terms of social welfare, because both mechanisms always build and they are both budget-balanced. It is never a good idea to set x to be strictly higher than C : if the non-excluded agents' total valuation is at least C , regardless of the excluded agent's type, the optimal decision is to build.

For any $x \in [0, C]$, mechanism x is (ex post) individually rational. Consider an arbitrary agent i . If agent i reports C/n , then he is never pivotal, so he does not pay any VCG payment excluded or not. If the decision is to build, then his utility is θ_i plus the redistribution he received from the others, which is at least 0. If the decision is not to build, then his utility is C/n plus the redistribution he received from the others, which

⁶ *Faltings* is also (ex post) individually rational in all settings where the VCG mechanism is (ex post) individually rational.

is also at least 0. That is, every agent can guarantee a non-negative utility by reporting C/n . Combining this with the fact that the mechanism is incentive compatible, we can conclude that it is individually rational.

Since the mechanism is always budget-balanced, for the purpose of maximizing social welfare, we can ignore payments when optimizing over x . Thus, for this purpose, we can simplify mechanism x to:

- We exclude one agent uniformly at random.
- If the non-excluded agents' total valuation is at least x , then we build. Otherwise, we do not build.

Theorem 1. *For any $x \in [0, C]$, mechanism x is at most $\frac{n-1}{n}$ -competitive.*

Proof. Mechanism 0 always builds. Consider the type profile $(0, 0, \dots, 0)$. Under mechanism 0, the agents' total utility is 0. The agents' maximum possible total utility $\max\{C, \theta_N\} = \max\{C, 0\} = C$. Hence, mechanism 0 is at most 0-competitive.

Consider $x > 0$. Consider the type profile $(U, 0, \dots, 0)$, where U is a number larger than C . Under mechanism x , when the agent reporting U is excluded, the decision is not to build (the agents' total utility is C). Otherwise, the decision is to build (the agents' total utility is U). The agents' expected total utility is $\frac{1}{n}C + \frac{n-1}{n}U$. The agents' maximum possible total utility is U . $\lim_{U \rightarrow \infty} \frac{\frac{1}{n}C + \frac{n-1}{n}U}{U} = \frac{n-1}{n}$. Hence, mechanism x is at most $\frac{n-1}{n}$ -competitive. \square

Theorem 2. *Mechanism C is exactly $\frac{n-1}{n}$ -competitive.*

Proof. Under mechanism C , if agent i is excluded, then the agents' total utility is at least $\max\{C, \sum_{j \neq i} \theta_j\}$. Averaging over all i , the agents' expected total utility is then at least

$$\frac{1}{n} \sum_{i=1}^n \max\{C, \sum_{j \neq i} \theta_j\}.$$

The above expression is no less than

$$\frac{1}{n} \max\{nC, \sum_{i=1}^n \sum_{j \neq i} \theta_j\} = \max\{C, \frac{n-1}{n} \theta_N\}.$$

This is always greater than or equal to $\frac{n-1}{n}$ times $\max\{C, \theta_N\}$. That is, mechanism C is exactly $\frac{n-1}{n}$ -competitive (Theorem 1 has shown that it is at most $\frac{n-1}{n}$ -competitive). \square

Theorem 3. *For any $x \in [\frac{n-2}{n-1}C, C)$, mechanism x is also exactly $\frac{n-1}{n}$ -competitive.*

Proof. As a result of Theorem 1, we only need to prove that for any $x \in [\frac{n-2}{n-1}C, C)$, mechanism x is at least $\frac{n-1}{n}$ -competitive.

For all type profiles with $\theta_N \geq C$, the correct (optimal) decision is to build. That is, for these type profiles, mechanism x is no worse than mechanism C , as mechanism x has a lower threshold for building.

Thus, we only need to prove that mechanism x is $\frac{n-1}{n}$ -competitive for all type profiles with $\theta_N < C$. For these type profiles, the correct decision is not to build. That is, if $x_1 \leq x_2$, then for these type profiles, mechanism x_2 is no worse than mechanism x_1 , as mechanism x_2 has a higher threshold for building.

Therefore, we only need to prove that mechanism $\frac{n-2}{n-1}C$ is $\frac{n-1}{n}$ -competitive for all type profiles with $\theta_N < C$. In other words, we only need to prove that under mechanism $\frac{n-2}{n-1}C$, the agents' expected total utility is at least $\frac{n-1}{n}C$ for all type profiles with $\theta_N < C$.

There are three cases:

1. If under mechanism $\frac{n-2}{n-1}C$, the decision is to build with probability 1, then we have for all i , $\sum_{j \neq i} \theta_j = \theta_N - \theta_i \geq \frac{n-2}{n-1}C$. That is, $\sum_{i=1}^n (\theta_N - \theta_i) \geq \sum_{i=1}^n (\frac{n-2}{n-1}C)$. Rearranging, $(n-1)\theta_N \geq \frac{n(n-2)}{n-1}C$. Therefore, we have that the agents' total utility θ_N is at least $\frac{n^2-2n}{n^2-2n+1}C \geq \frac{n-1}{n}C$ (recall that $n \geq 3$).
2. If under mechanism $\frac{n-2}{n-1}C$, the decision is to build with probability $\frac{1}{n} \leq p \leq \frac{n-1}{n}$, then the agents' expected total utility is $p\theta_N + (1-p)C$. This expression is decreasing in p , and increasing in θ_N . It is minimized when $p = \frac{n-1}{n}$ and $\theta_N = \frac{n-2}{n-1}C$ ($\theta_N \geq \frac{n-2}{n-1}C$ because there exists i such that $\theta_N - \theta_i \geq \frac{n-2}{n-1}C$). That is, the agents' expected total utility is minimized under type profile $(\frac{n-2}{n-1}C, 0, 0, \dots, 0)$. For this type profile, the agents' expected total utility is $\frac{1}{n}C + \frac{n-1}{n} \frac{n-2}{n-1}C = \frac{n-1}{n}C$.
3. If under mechanism $\frac{n-2}{n-1}C$, the decision is to build with probability 0, then this mechanism is always making the correct decision. The agents' total utility is C .

□

Theorem 4. For any $x \in [0, \frac{n-2}{n-1}C)$, mechanism x is strictly less than $\frac{n-1}{n}$ -competitive.

Proof. We have already shown that mechanism 0 is at most 0-competitive in the proof of Theorem 1.

For $x > 0$, consider the type profile $(x, 0, 0, \dots, 0)$. Under mechanism x , when the agent reporting x is excluded, the decision is not to build, and the agents' total utility is C . When some other agent is excluded, the decision is to build, and the agents' total utility is x . The agents' expected total utility is $\frac{n-1}{n}x + \frac{1}{n}C$. The agents' maximum possible total utility is C . The ratio equals $\frac{(n-1)x}{nC} + \frac{1}{n} < \frac{(n-1)\frac{n-2}{n-1}C}{nC} + \frac{1}{n} = \frac{n-1}{n}$. □

As a summary, we have shown that mechanism x is optimal if and only if $x \in [\frac{n-2}{n-1}C, C]$. Next we consider mixtures of mechanisms with different parameters.

Definition 1. Mechanism OptMix:

- With probability $\frac{1}{n+1}$, we run mechanism 0 (always build);
- With probability $\frac{n}{n+1}$, we run mechanism C .

Theorem 5. OptMix is exactly $\frac{n}{n+1}$ -competitive.

We note that OptMix is more competitive than any individual (non-mixture) mechanism x .

Proof. If $\theta_N < C$, then mechanism C never builds. That is, if $\theta_N < C$, then the agents' expected total utility under $OptMix$ is $\frac{1}{n+1}\theta_N + \frac{n}{n+1}C \geq \frac{n}{n+1}C = \frac{n}{n+1} \max\{C, \theta_N\}$.

We have that mechanism C is $\frac{n-1}{n}$ -competitive, so if $\theta_N \geq C$, then the agents' expected total utility under mechanism C is at least $\frac{n-1}{n}\theta_N$. Under $OptMix$, the agents' expected total utility is then at least $\frac{1}{n+1}\theta_N + \frac{n}{n+1}\frac{n-1}{n}\theta_N = \frac{n}{n+1}\theta_N = \frac{n}{n+1} \max\{C, \theta_N\}$.

The above shows that $OptMix$ is at least $\frac{n}{n+1}$ -competitive. Let us consider the type profile $(0, 0, \dots, 0)$. For this type profile, under $OptMix$, the agents' expected total utility is exactly $\frac{n}{n+1}C = \frac{n}{n+1} \max\{C, \theta_N\}$. Hence, $OptMix$ is exactly $\frac{n}{n+1}$ -competitive. \square

Let Mix be an arbitrary mixture of mechanisms with different parameters. Let I be an interval that is a subset of $[0, C]$. We use $P(Mix \in I)$ to denote the probability that a mechanism with parameter $x \in I$ is used. $P(Mix \in [0, C]) = 1$. For $OptMix$, we have $P(OptMix \in [0, 0]) = \frac{1}{n+1}$ and $P(OptMix \in [C, C]) = \frac{n}{n+1}$. We will prove that Mix is at most $\frac{n}{n+1}$ -competitive. That is, $OptMix$ is the most competitive among all mixtures of mechanisms with different parameters.⁷

Theorem 6. *OptMix is the most competitive among all mixtures of mechanisms with different parameters.*

Proof. If $P(Mix \in [0, 0]) < \frac{1}{n+1}$, then let us consider the type profile $(U, 0, \dots, 0)$, where U is larger than C . When the agent reporting U is excluded (which happens with probability $\frac{1}{n}$), the non-excluded agents' types are all zeros, which means that the probability to build (when the agent reporting U is excluded) is equal to $P(Mix \in [0, 0])$. That is, overall, for this type profile, the probability \bar{p} of not building is at least $\frac{1}{n}(1 - P(Mix \in [0, 0]))$, which is strictly larger than $\frac{1}{n}(1 - \frac{1}{n+1}) = \frac{1}{n+1}$. The agents' expected total utility is $(1 - \bar{p})U + \bar{p}C$. The agents' maximum possible total utility is U . We have $\lim_{U \rightarrow \infty} \frac{(1-\bar{p})U + \bar{p}C}{U} = 1 - \bar{p}$. That is, if $P(Mix \in [0, 0]) < \frac{1}{n+1}$, then Mix is at most $\frac{n}{n+1}$ -competitive. Therefore, if Mix is to be no less competitive than $OptMix$, then we must have $P(Mix \in [0, 0]) \geq \frac{1}{n+1}$.

If $P(Mix \in [0, 0]) \geq \frac{1}{n+1}$, then let us consider the type profile $(0, 0, \dots, 0)$. The probability \bar{p} of not building is most $1 - P(Mix \in [0, 0])$. It follows that $\bar{p} \leq \frac{n}{n+1}$. The agents' expected total utility is $(1 - \bar{p})0 + \bar{p}C$. The agents' maximum possible total utility is C . The ratio equals \bar{p} , which is at most $\frac{n}{n+1}$, and it follows that Mix is at most $\frac{n}{n+1}$ -competitive. \square

So far, we have identified many competitive randomized mechanisms. Another natural question to ask is whether there exist competitive deterministic mechanisms. The answer is yes: the VCG mechanism is $\frac{1}{n}$ -competitive.

⁷ $OptMix$ is not the unique optimum. Consider a modified version of $OptMix$ under which we run mechanism 0 with probability $\frac{1}{n+1}$ and run mechanism $C - \epsilon$ with probability $\frac{n}{n+1}$ (ϵ is a small positive number). When $\theta_N \geq C$, modified $OptMix$ is no worse than $OptMix$, since the optimal decision is to build, and modified $OptMix$ has a lower threshold for building. When $\theta_N < C - \epsilon$, modified $OptMix$ is the same as $OptMix$. When $\theta_N \in [C - \epsilon, C)$, the optimal decision is not to build. The maximum efficiency is C . Under any budget-balanced mechanism (including modified $OptMix$), the agents' expected total utility is between $C - \epsilon$ and C , thus at least $C - \epsilon$. When ϵ is small enough, we have $\frac{C-\epsilon}{C} \geq \frac{n}{n+1}$.

Theorem 7 (Moulin, private communication). *The VCG mechanism is exactly $\frac{1}{n}$ -competitive for the public project problem.*

To illustrate how poor of a ratio the $1/n$ achieved by VCG is, we now give a very simple mechanism that also obtains this ratio.

Definition 2. *Mechanism Vote-to-Build: Let every agent vote whether to build or not. If there is at least one vote toward building, then we build. Otherwise, we do not build.*

If an agent's valuation is at least C/n , then his dominant strategy is to vote toward building. If an agent's valuation is less than C/n , then his dominant strategy is to vote toward not building. *Vote-to-Build* is (ex post) individually rational and budget-balanced (there are no payments involved).

Theorem 8. *Vote-to-Build is exactly $\frac{1}{n}$ -competitive.*

Proof. If the decision is to build, then there exists i with $\theta_i \geq C/n$. That is, we have $\theta_N \geq C/n$. The ratio $\frac{\theta_N}{\max\{C, \theta_N\}}$ is at least $\frac{1}{n}$, and it reaches $\frac{1}{n}$ when $\theta_N = C/n$ (corresponding to the type profile $(C/n, 0, 0, \dots, 0)$).

If the decision is not to build, then there is no i with $\theta_i \geq C/n$. It follows that $\theta_N \leq C$. The ratio is then $\frac{C}{\max\{C, \theta_N\}} = \frac{C}{C} = 1$. \square

4 General Domains

In Section 3, we showed that *Faltings* is at least $\frac{n-1}{n}$ -competitive for the public project problem. Here we show that it remains true for general mechanism design problems, as long as the agents' valuation functions satisfy the following assumption.⁸

Assumption 1 *The valuations of agents are non-negative for all outcomes: $v_i(k, \theta_i) \geq 0$ for all θ_i, k .*

Theorem 9. *Faltings is at least $\frac{n-1}{n}$ -competitive, as long as the agents' valuations satisfy Assumption 1.*

Proof. We begin with observing a few relationships between the value of the efficient outcome when all agents are present and when one agent is excluded. Under Assumption 1, making agent i accept a decision made without him does not decrease the value of that decision.⁹ In particular, this applies to the efficient outcome for agents $j \neq i$:

$$\sum_j v_j(k^*(\theta_{-i}), \theta_j) \geq \sum_{j \neq i} v_j(k^*(\theta_{-i}), \theta_j).$$

On the other hand, the total value of agents $j \neq i$ under the outcome efficient for them is at least as high as their total value under the outcome efficient when all agents are present.

⁸ The assumption places restrictions only on the valuation function and is independent of the mechanism. This is in contrast to the individual rationality property, which requires the *utility* of each agent participating in the mechanism to be above his outside value.

⁹ The valuation function in Equation 1 satisfies this property.

$$\sum_{j \neq i} v_j(k^*(\theta_{-i}), \theta_j) \geq \sum_{j \neq i} v_j(k^*(\theta), \theta_j).$$

Combining the two inequalities and summing over all agents, we get

$$\begin{aligned} \sum_i \sum_j v_j(k^*(\theta_{-i}), \theta_j) &\geq \sum_i \sum_{j \neq i} v_j(k^*(\theta_{-i}), \theta_j) \\ &\geq \sum_i \sum_{j \neq i} v_j(k^*(\theta), \theta_j) = (n-1) \sum_i v_i(k^*(\theta), \theta_i). \end{aligned}$$

Dividing by n and focusing on the first and last expressions, we have

$$\frac{1}{n} \sum_i \sum_j v_j(k^*(\theta_{-i}), \theta_j) \geq \frac{n-1}{n} \sum_i v_i(k^*(\theta), \theta_i).$$

The expression on the left-hand side is the expected value of the outcome when the decision is made efficiently after one agent is excluded uniformly at random. The inequality implies that the expected value of the decision under *Faltings* is at least $\frac{n-1}{n}$ of the maximum efficiency. \square

Faltings results in a rather unequal treatment of the excluded agent relative to the other agents. In settings where the VCG mechanism collects a lot of revenue, the agents would be envious of the excluded agent.

We next study a more fair payment scheme where each agent pays his expected VCG payment and receives part of his own residual claimant rebate. We call the resulting mechanism *FaltingsFair*. This scheme had been proposed previously by Faltings [6, 8]. We derived it independently with formal proofs. The result on the competitive ratio is novel. We discuss this in more detail at the end of this section.

- Exclude an agent a uniformly at random and compute the efficient allocation.
- Collect from each agent i (including a) the payment

$$t_i(\theta) = \frac{1}{n} \sum_{j \neq i} t_i^{\text{vcg}}(\theta_{-j}) - \frac{1}{n} \sum_{j \neq i} t_j^{\text{vcg}}(\theta_{-i}) \quad (2)$$

Expanding each term of the payment, we can rewrite it as follows.

$$\begin{aligned} t_i^{\text{vcg}}(\theta_{-j}) &= \sum_{a \neq i, j} v_a(k^*(\theta_{-i, j}), \theta_a) - \sum_{a \neq i, j} v_a(k^*(\theta_{-j}), \theta_a) \\ t_j^{\text{vcg}}(\theta_{-i}) &= \sum_{a \neq i, j} v_a(k^*(\theta_{-i, j}), \theta_a) - \sum_{a \neq i, j} v_a(k^*(\theta_{-i}), \theta_a) \\ t_i^{\text{vcg}}(\theta_{-j}) - t_j^{\text{vcg}}(\theta_{-i}) &= \sum_{a \neq i, j} (v_a(k^*(\theta_{-i}), \theta_a) - v_a(k^*(\theta_{-j}), \theta_a)) \\ t_i(\theta) &= \frac{1}{n} \sum_{j \neq i} \sum_{a \neq i, j} (v_a(k^*(\theta_{-i}), \theta_a) - v_a(k^*(\theta_{-j}), \theta_a)) \end{aligned}$$

Theorem 10. *FaltingsFair is incentive compatible in expectation, budget-balanced, and for valuations satisfying Assumption 1, $\frac{n-1}{n}$ -competitive.*

Proof. First we prove incentive compatibility. Denoting *FaltingsFair*'s allocation function that chooses a residual claimant uniformly at random with k^{rc} , agent i 's utility

$$\begin{aligned}
u_i(k^{\text{rc}}(\theta), t_i, \theta_i) &= \left(\frac{1}{n} \sum_j v_i(k^*(\theta_{-j}), \theta_i) \right) - t_i(\theta) \\
&= \frac{1}{n} \sum_j v_i(k^*(\theta_{-j}), \theta_i) - \frac{1}{n} \sum_{j \neq i} \sum_{a \neq i, j} (v_a(k^*(\theta_{-i}), \theta_a) - v_a(k^*(\theta_{-j}), \theta_a)) \\
&= \frac{1}{n} v_i(k^*(\theta_{-i}), \theta_i) + \frac{1}{n} \sum_{j \neq i} \left(\sum_{a \neq j} v_a(k^*(\theta_{-j}), \theta_a) - \sum_{a \neq i, j} v_a(k^*(\theta_{-i}), \theta_a) \right).
\end{aligned}$$

Removing the terms that agent i does not control with his report, we are left with $\frac{1}{n} \sum_{j \neq i} \sum_{a \neq j} v_a(k^*(\theta_{-j}), \theta_a)$. This expression is maximized when agent i reports the true value θ_i as by the definition of $k^*(\theta_{-j})$

$$\sum_{a \neq j} v_a(k^*(\theta_{-j}), \theta_a) \geq \sum_{a \neq j} v_a(k', \theta_a) \quad \forall j, k' \in K.$$

Therefore, incentive compatibility holds.

Next we show budget balance ($\sum_i t_i = 0$).

$$\begin{aligned}
\sum_i \frac{1}{n} \left(\sum_{j \neq i} t_i^{\text{vcg}}(\theta_{-j}) - \sum_{j \neq i} t_j^{\text{vcg}}(\theta_{-i}) \right) &= 0 \\
\sum_i \sum_{j \neq i} t_i^{\text{vcg}}(\theta_{-j}) &= \sum_i \sum_{j \neq i} t_j^{\text{vcg}}(\theta_{-i})
\end{aligned}$$

The equality follows from the simple identity $\sum_i \sum_{j \neq i} a_{ij} = \sum_i \sum_{j \neq i} a_{ji}$.

Finally, the allocation function is the same as before, thus, *FaltingsFair* has the same competitive ratio as *Faltings*. \square

Unlike *Faltings*, *FaltingsFair* is incentive compatible only in expectation with respect to the random outcome function $k(\theta)$. This means that an agent has no incentive to misreport his value before the outcome is chosen,¹⁰ but once the outcome is known, the agent may regret not reporting a different value. Incentive compatibility in expectation is a natural concept for randomized mechanisms as the reporting of values must occur before the outcome is selected.

Our next theorem deals with individual rationality. [7] showed that *Faltings* is (ex post) individually rational in settings where the VCG mechanism is (ex post) individually rational. This is actually the case for valuations satisfying Assumption 1. That is, for valuations that satisfy Assumption 1, *Faltings* is (ex post) individually rational. Similar to the case of incentive compatibility, unlike *Faltings*, *FaltingsFair* is individually rational only in expectation with respect to the random outcome function $k(\theta)$.

Theorem 11. *For valuations satisfying Assumption 1, FaltingsFair is individually rational in expectation.*

We now take a closer look at the payment function in Equation 2. The first term is the expected VCG payment in the market with one agent excluded uniformly at random. The second term produces a rebate equal to $\frac{1}{n}$ of the total VCG payments realized without the agent in the market. This rebate has been considered before with the goal of redistributing the VCG surplus in [3, 4].

$$h_i^{\text{rc}}(\theta_{-i}) = \frac{1}{n} \sum_{j \neq i} t_j^{\text{vcg}}(\theta_{-i}) \quad (3)$$

¹⁰ Note that unlike the Bayesian incentive compatible “expected externality mechanism” (or dAGVA), our mechanism is dominant-strategy incentive compatible and we have no prior over agents’ types.

This rebate, however, may exceed the total VCG revenue resulting in a deficit in some models. In fact, as Cavallo argues in [4], the no-deficit property requires the redistribution to sometimes be smaller than the amount above. Specifically, one can compute the smallest total VCG payment collected from the agents over all values agent i might have. It is this amount that should be used in the redistribution to agent i :

$$h_i^{\min}(\theta_{-i}) = \frac{1}{n} \min_{\theta'_i} \sum_j t_j^{\text{vCG}}(\theta'_i, \theta_{-i}) \quad (4)$$

It is easy to see that in the public good setting, the above rebate (Equation 4) is always zero.¹¹ Thus, rebates of this form are not helpful in efficient mechanisms in models like those involving public goods. In contrast, in *FaltingsFair*, the rebate (3) results in full budget balance in any model.

The payment rule in Equation 2 was previously proposed by Boi Faltings in a patent [8] and an unpublished paper [6]. There Faltings provides an equivalent definition of the payment rule: instead of considering the rebate function explicitly, the rule directs each agent i to pay $\frac{1}{n} t_i^{\text{vCG}}(\theta_{-j})$ to each agent j . Notice that budget balance follows immediately from this definition. To see that the definition in fact defines the payment rule in Equation 2, notice that the first summation corresponds to the payments agents $j \neq i$ make to agent i and the second summation corresponds to the payments agent i makes to agents $j \neq i$. These definitions provide different interpretations of the mechanism: Faltings views it as the average of the budget-balanced *Faltings* mechanism, while we make explicit the connections to a redistribution function previously considered in the literature.

5 Discussion

We studied randomized mechanisms that are fully budget-balanced and aimed to maximize the expected efficiency, which under budget balance coincides with the expected social welfare. The expected welfare loss of our generally applicable mechanism is only $\frac{1}{n}$ in the worst case leaving little room for improvement. However, whether or not this loss can be reduced with a different randomized mechanism (budget-balanced, or not) remains an open question.

Note that full efficiency is impossible in randomized mechanisms.¹² Thus, the goal of minimizing budget imbalance in an efficient mechanism is not meaningful in this context. However, the question of minimizing budget imbalance in deterministic mechanisms for public good remains open.

Finally, we note that for public project problems, our definition of the value of the efficient outcome adds C to the standard definition of $\max(\theta_N - C, 0) = \max(\theta_N, C) - C$ (see, e.g., [13]). Our results can be interpreted under the standard definition: being

¹¹ If $\theta_N - \theta_i < \frac{n-1}{n}C$, then when $\theta_i = 0$, no agent is pivotal and the total VCG payment is 0; If $\theta_N - \theta_i \geq \frac{n-1}{n}C$, then when $\theta_i = C$, no agent is pivotal and the total VCG payment is also 0.

¹² At least, this is the case for *significantly* randomized mechanisms that do more than just break ties randomly.

r -competitive means we guarantee the welfare of $r(C + \max(\theta_N - C, 0))$. We cannot guarantee the welfare of $r \max(\theta_N - C, 0)$ as Assumption 1 does not hold for the standard valuation function

$$v_i(k(\theta), \theta_i) = \begin{cases} \theta_i - \frac{C}{n} & \text{if } k(\theta) = 1 \\ 0 & \text{otherwise} \end{cases}$$

We are indebted to Hervé Moulin for the idea of using the alternative metric.

The definition of valuations above ensures that the cost C is covered (each agent contributes $\frac{C}{n}$) as long as the sum of the agents' payments is non-negative. The valuation function that we use shifts the standard one by $\frac{C}{n}$. However, in both cases, the cost C is covered: i.e., non-negative total payments result in weak budget balance.

References

1. K.R. Apt, V. Conitzer, M. Guo, and E. Markakis. Welfare undominated groves mechanisms. In *WINE*, pages 426–437, 2008.
2. K.R. Apt and A. Estévez-Fernández. Sequential pivotal mechanisms for public project problems. In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory, SAGT '09*, pages 85–96, Berlin, Heidelberg, 2009. Springer-Verlag.
3. M.J. Bailey. The demand revealing process: To distribute the surplus. *Public Choice*, 91(2):107–26, April 1997.
4. R. Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *AAMAS'06*, Hakodate, Japan, 2006.
5. G. de Clippel, V. Naroditskiy, and A. Greenwald. Destroy to save. In *EC'09*, 2009.
6. B. Faltings. A budget-balanced, incentive-compatible scheme for social choice. Unpublished, 2004.
7. B. Faltings. A budget-balanced, incentive-compatible scheme for social choice. In Peyman Faratin and Juan A. Rodríguez-Aguilar, editors, *Agent-Mediated Electronic Commerce VI*, volume 3435, pages 30–43. Springer, 2005.
8. B. Faltings. Social choice determination systems and methods, 2011. United States Patent.
9. S. Gujar and Y. Narahari. Redistribution mechanisms for assignment of heterogeneous objects. *Journal of Artificial Intelligence Research*, 41:131–154, 2011.
10. M. Guo and V. Conitzer. Better redistribution with inefficient allocation in multi-unit auctions with unit demand. In *EC'08*, pages 210–219, 2008.
11. M. Guo and V. Conitzer. Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games and Economic Behavior*, 67(1):69 – 98, 2009.
12. M. Guo and V. Conitzer. Optimal-in-expectation redistribution mechanisms. *Artif. Intell.*, 174(5-6):363–381, 2010.
13. A. Mas-Colell, M.D. Whinston, and J.R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
14. H. Moulin. Almost budget-balanced VCG mechanisms to assign multiple objects. *Journal of Economic Theory*, 144(1):96–119, 2009.